



(11) **EP 1 400 900 B1**

(12) **EUROPEAN PATENT SPECIFICATION**

(45) Date of publication and mention
of the grant of the patent:
28.12.2016 Bulletin 2016/52

(51) Int Cl.:
G06F 12/08 ^(2006.01) **G06F 12/12** ^(2006.01)
G06F 17/30 ^(2006.01)

(21) Application number: **03255798.5**

(22) Date of filing: **16.09.2003**

(54) **Selective cache admission into a web cache**

Selektiver cachespeicherzugang zu einen WEB Cachespeicher

Admission sélective en anté mémoire pour une anté mémoire WEB

(84) Designated Contracting States:
DE DK FI FR GB

(30) Priority: **17.09.2002 US 244449**

(43) Date of publication of application:
24.03.2004 Bulletin 2004/13

(73) Proprietor: **Ill Holdings 3, LLC**
Wilmington, DE 19808 (US)

(72) Inventors:
• **Dixit, Sudhir**
Weston, MA 02493-1151 (US)

• **Wu, Tau**
Woburn, MA 01801 (US)

(74) Representative: **Carpmaels & Ransford LLP**
One Southampton Row
London WC1B 5HA (GB)

(56) References cited:
• **KURCEWICZ M ET AL: "A filtering algorithm for
Web caches" COMPUTER NETWORKS AND ISDN
SYSTEMS, NORTH HOLLAND PUBLISHING.
AMSTERDAM, NL, vol. 30, no. 22-23, 25 November
1998 (1998-11-25), pages 2203-2209,
XP004152172 ISSN: 0169-7552**

Note: Within nine months of the publication of the mention of the grant of the European patent in the European Patent Bulletin, any person may give notice to the European Patent Office of opposition to that patent, in accordance with the Implementing Regulations. Notice of opposition shall not be deemed to have been filed until the opposition fee has been paid. (Art. 99(1) European Patent Convention).

EP 1 400 900 B1

Description

FIELD OF THE INVENTION

[0001] This invention relates to selective admission of data into memory. In particular, this invention relates to systems and methods for selectively admitting objects into, e.g., a web cache.

BACKGROUND OF THE INVENTION

[0002] The Internet, and in particular, the World Wide Web (WWW or web), is becoming an integral part of modern life. Unfortunately, the growth of the web places ever-increasing demands on the network backbone and other facilities that form the web. Web traffic has been growing at a much faster pace than available bandwidth, often causing substantial latency between user request for content and user receipt of that content. In many cases, this latency results from network congestion caused by numerous requests for transmission of the same content. Such activities can overload (and in some cases, disable) web servers and other network facilities. At a minimum, multiple requests for the same material from a web server increase delays experienced by web users.

[0003] Web caching offers potential relief to overloaded networks. As is known in the art, web caching is a technique of storing popular web content at, and providing that stored content to end users from, locations in addition to the web servers that initially provide that content. By making copies of web pages and other content available from alternate locations, the load upon the origin servers that initially provide the content is reduced, substantially reducing latency. Web caching also helps transfer load from the Internet backbone to smaller networks. By storing frequently requested web content at one or more web cache servers located at network edge(s), future local requests for that content can be served from the web cache(s) instead of repeatedly obtaining content from the origin servers. This reduces Internet traffic, and may also reduce load upon Wide Area Networks (WANs) and other networks that are linked by (or to) the Internet. Load on origin web servers is reduced because those origin servers service fewer requests.

[0004] Web caches may be deployed in numerous and varied configurations. FIGS. 1 and 2 represent only a few examples. Both FIG. 1 and FIG. 2 illustrate deployment scenarios in which the existence of the web server is not apparent to the end user/client. It is possible that no manual or automatic configuration of client web browser software is needed to access the web cache (although the web cache may only serve users within a specific network), and the user may perceive no difference between content requests served by a web cache vs. content requests served by an origin server. FIG. 1 illustrates a typical web cache deployed at a network edge. In this scenario, clients on a local network send HTTP (Hypertext Transfer Protocol) requests to origin servers on the

Internet. These requests may be forwarded by a local network router within the local network to a switch. That switch may have Layer 4 (transport layer) or Layer 7 (application layer) capability, and thus be able to identify HTTP traffic.

[0005] For example, a Layer 4 switch might identify HTTP traffic by checking the TCP (Transmission Control Protocol) port number of incoming IP (Internet Protocol) packets. If the destination port number is 80 (default HTTP server port number), the packet is forwarded to the cache. Otherwise, the packet could be forwarded to the WAN Router. The cache then intercepts the TCP connection from the client and obtains the URL (Universal Resource Locator) for the desired Web pages or other content. A Layer 7 switch (also known as a content switch or web switch) may replace the Layer 4 switch to provide additional functionality. For example, TCP connections from clients may be intercepted by a Layer 7 switch instead of the cache, and the Layer 7 switch might make routing decisions based on the URL. In either event, a switch identifies HTTP traffic and forwards that traffic to the cache. If the content requested by the client is stored in the cache, that content is provided to the client from the cache. Otherwise, the cache fetches the content from an origin server or other location, and serves the content to the requesting client.

[0006] FIG. 2 illustrates a typical reverse proxy scenario where web caches are used to relieve the load upon web servers. Incoming requests are intercepted by a Layer 7 switch. Based on how the reverse proxy is configured, either a cache or server is selected to serve the request. For example, frequently changing content may generally be served by a web server, and relatively unchanging content served by a web cache. Because the cost of a web cache is typically much lower than the cost of a web server, deploying web caches to serve popular static content provides an economic and scalable server farm solution.

[0007] In both scenarios shown by FIGS. 1 and 2, as well as in other scenarios, web caching improves user experience and relieves load on origin servers. If deployed at a network edge, web caching can also provide substantial cost savings in terms of backbone bandwidth. Other aspects of web caching may undercut these benefits, however. In a steady state, a web cache optimally operates at full (or near-full) storage capacity. Accordingly, before a new object may be stored in the cache, one or more old objects must be evicted from the cache. Various cache replacement policies have been developed to optimize the eviction process based on measurements such as maximizing Hit Ratio (ratio of requests served by cache to all requests received by cache) or minimizing user perceived latency.

[0008] However, web caching has unique characteristics that must be addressed. Unlike caching in a memory hierarchy using fixed-size blocks, web caching must accommodate web objects of widely varying size. Moreover, an overloaded or improperly configured web cache

may itself become a network bottleneck and increase latency rather than decrease latency. Typically, web caches store actual content in hard disk drives or in other storage devices that have relatively slow moving mechanical parts. These devices support a relatively limited number of operations per second; these operations include storing new objects as well as accessing stored objects. In other words, time spent storing new objects is generally at the expense of time that might be used to access previously stored objects. Unless the number of disk (or other device) I/O operations are controlled in some manner, the throughput of the cache is not optimized.

[0009] To date, there have been limited solutions to these problems. As one example, a Layer 7 switch can be deployed as in FIG. 1, and configured to bypass the cache when the cache becomes overloaded. This approach increases traffic on the network backbone and does not address the underlying cause of cache overload. Multiple hard drives (or even multiple caches) can be deployed in parallel so as to improve total cache throughput, but this solution requires increased hardware investment.

[0010] Accordingly, there remains a need for improved methods and systems of managing web cache storage.

[0011] One prior approach is described in Kurcewicz M et al: "A filtering algorithm for web caches" Computer Networks and ISDN Systems, North Holland Publishing, Amsterdam, NL, Vol. No. 30, no. 22-23, 25 November 1998 pp 2203-2209, XP004152172, ISSN:0169-7552. This describes a filtering algorithm that caches web objects selectively depending on whether they have been requested by more than one client.

SUMMARY OF THE INVENTION

[0012] The present invention is set out in the claims hereinafter. According to the invention, if an object is not stored in the memory device, but has previously been requested a designated number of times or is of a size that satisfies a size criterion, the object is stored. In one embodiment, the object is admitted upon a second request regardless of its size, and is admitted on a first request if it is smaller than the average size of objects currently stored in the memory. To make room for new objects, other objects are evicted from the memory on, e.g., a Least Recently Used (LRU) basis. The invention could be implemented on existing web caches, on distributed web caches, and in client-side web caching. The invention could further be implemented in connection with storing data that may be unrelated to Internet content.

BRIEF DESCRIPTION OF THE DRAWINGS

[0013]

FIG. 1 is an example of a web cache deployed within

a network;

FIG. 2 is another example of a web cache deployed within a network;

FIG. 3 is a schematic drawing of a web cache servicing a request for an object;

FIG. 4 is another schematic drawing of a web cache servicing a request for an object;

FIG. 5 is a schematic drawing of an object being admitted into a web cache;

FIG. 6 is a flow chart showing one example of a method according to the invention;

FIG. 7 is a graph showing a comparison of Hit Ratio for the invention compared to web cache management using only LRU eviction;

FIG. 8 is a graph showing a comparison of insertion operations per request for the invention compared to web cache management using only LRU eviction;

FIG. 9 is a graph showing a comparison of eviction operations per request for the invention compared to web cache management using only LRU eviction; and

FIG. 10 is a graph showing a comparison of disk operations per request for the invention compared to web cache management using only LRU eviction.

DETAILED DESCRIPTION OF THE INVENTION

[0014] The present invention implements admission control to selectively admit data into a memory. As used herein, and unless otherwise specified, "memory" includes both non-volatile data storage (e.g., hard disk drives, optical drives, etc.) and volatile memory (e.g., RAM). The invention may advantageously be implemented in a web cache, and will be described using a web cache as an example. The invention is not limited to such implementation, however. The invention may be used to improve memory management in client side caching, or in general data caching that may be unrelated to Internet content. In that vein, use of the word "caching" in this description to indicate storage of an object in a web cache should not be construed as limiting the invention to data storage environments previously referred to as caches.

[0015] One metric often used to evaluate performance of web caches is the Hit Ratio:

55

$$\text{Hit Ratio} = \frac{\text{requests successfully served}}{\text{total requests}}$$

Caching memory systems in contexts other than web caching often achieve Hit Ratios exceeding 95%. In network edge web cache deployment scenarios, however, approximately half of all requests appear only once. By implication, a web cache Hit Ratio generally will not exceed 50% to 60% under the best of conditions. Evicting an object from a web cache to accommodate a new object is thus more likely to exceed the benefit of bringing the new object into the web cache. Moreover, in the case of a web cache using a hard disk drive or other storage device having moving mechanical parts, bringing a new object (which may be in the half of objects not requested more than once) into the web cache may require time-consuming operations (e.g., movement of a read/write arm).

[0016] Another measurement of web cache efficiency is throughput, or the amount of content served by the web cache over time. For a web cache using a hard disk drive or other storage device having moving mechanical parts, throughput can be improved by reducing the total number of input/output (I/O) operations on the disk or other device; time spent writing new objects into a web cache detracts from time available to serve requests with objects already stored in the web cache. Without admission control, simply improving Hit Ratio does not enhance throughput in such systems. This is shown by observing that without admission control, a web cache responds to each incoming request by either serving a cached object (e.g., a disk read operation) or caching a new object upon retrieval from another server (e.g., a disk write operation). Because there are only two possibilities for each request of a web cache without admission control, the following equation generally governs:

$$H + I = 1,$$

where H is the Hit Ratio, and I is the statistical number of insertion operations per request. For example, if Hit Ratio is 45% (45 out of every 100 requests are served from the cache), I is 55% (55 requests out of every 100 cause a new object to be written into the cache). Other operations, such as evicting an object from the cache, are typically performed in main memory (e.g., RAM) and do not require a disk operation.

[0017] If admission control is implemented, the total number of I/O operations for a disk or other device can be reduced. If the requested object is in the cache, a read operation occurs when the object is served. If the requested object is not in the cache, however, it is not necessarily cached upon retrieval from another web server. If the cache storage spends less time writing objects into the cache, more time is available to serve requests. By appropriately controlling admission of new objects into the cache, both Hit Ratio and throughput may thus be improved.

[0018] In the context of web caches, object size is gen-

erally unrelated to the frequency with which the object is requested. If an object is not in a web cache and has not previously been requested from the web cache, an embodiment of the present invention admits the new object only if its size is smaller than the average size of currently cached objects. This reduces the number of evictions per new admitted object, as the new object is statistically smaller than the object(s) being evicted. This also improves Hit Ratio. Statistically, more objects can be cached by reducing average object sizes. If an object fails the admission control test on the first request, the object is then admitted upon a subsequent request. This allows the web cache to store large but popular objects. A candidate list may be maintained by the web cache to store the URLs of each first time request. Because only the URLs are stored, and not the actual objects, the candidate list can reside in a web cache server's main memory (e.g. volatile RAM); no disk I/O operation is necessary to add or remove candidate objects from the list.

[0019] FIGS. 3 through 5 schematically show operation of one embodiment of the invention. FIG. 3 shows web cache 100, which includes cache storage 105 and cache manager 110. Cache storage 105 can be a hard disk drive (or multiple hard disk drives) or some other device for storage of large amounts of data, but the invention could be implemented using any type of memory. Stored upon cache storage 105 are multiple objects 0001, 0002, 0003, etc. These objects may include text files, graphics files, video or audio files, applets, or any other type of data (or software or software/data combination) that might be requested over the Internet or other network. Cache manager 110 includes software that controls operation of web cache 100, and that may reside on (and control the operation of) a web cache. Cache manager 110 could alternately reside on (and control operation of) other computing devices having the necessary processor(s) and memory.

[0020] Through appropriate network interconnection(s), such as but not limited to those shown in FIGS. 1 and 2, cache manager 110 receives incoming requests for web objects. Cache manager 110 also controls retrieval of objects stored in cache storage 105, as well as writing new objects into storage 105. Cache manager 110 further maintains one or more data files to track objects stored in cache storage 105, as well as to track requests for objects. Those data files are shown schematically in FIGS. 3-5 as files 120 and 122. Two such files are shown for purposes of illustrating the operation of the invention. However, persons skilled in the art will appreciate that the types of information contained in those data files (explained below) need not be in a specific arrangement. In other words, the information could be stored in only one file or in more than two files. In one embodiment, files 120 and 122 may be maintained in the main memory (RAM) of a web cache, and thus accessible without requiring an I/O operation of cache storage 105 or of some other storage device having moving mechanical components.

[0021] As shown in FIG. 3, cache 100 receives a request for object *k*. Cache manager 110 then determines from file 120 that object *k* is stored in cache storage 105. A copy of object *k* is then retrieved from cache storage 105 and served in response to the request. File 120 includes various data about the objects in cache storage 105, such as the size of each object and the date and time when the object was last requested. Cache manager 110 also calculates an average size of the objects in cache storage 105. Upon serving a copy of object *k* in response to the request, cache manager 110 updates file 120 with the date and/or time of the request.

[0022] FIG. 4 illustrates a request for an object *x* that is not currently stored within cache storage 105. After determining that object *x* is not identified in file 120, web cache 100 retrieves object *x* from origin server 210. Object *x* is retrieved via one or more networks and/or the Internet, represented collectively as cloud 200. Upon obtaining object *x*, web cache 100 serves object *x* in response to the request. Cache manager 110 also determines whether object *x* should be stored in cache storage 105. Cache manager 110 first determines from file 122 whether object *x* has previously been requested. Because object *x* has not previously been requested, it is not stored in cache storage 105 at that point. Cache manager 110 then determines from file 120 (see FIG. 3) whether the size of object *x* is less than the average size of objects already stored in cache storage 105. Because the size of object *x* is not less than the average, cache manager 110 still does not store object *x* in cache storage 105. However, as shown in FIG. 4, file 122 is updated to indicate the time and/or date of the current request, and optionally, the number of prior requests. Persons skilled in the art will appreciate that these determinations need not occur in the described order. For example, cache manager 110 might first compare object *x* to a size criterion, and upon determining that it does not satisfy the criterion, determine if the object has previously been requested.

[0023] FIG. 5 illustrates a subsequent request for object *x*. Upon receiving the subsequent request, web cache 100 again retrieves object *x* from origin server 210, and serves object *x* in response to the request. Cache manager 110 then determines from file 122 that object *x* has previously been requested. Cache manager 110 then evicts object *y* from cache storage 105 to make room for object *x*. In one embodiment, eviction may occur by simple identification of the currently stored object(s) (from file 120) for which the currently occupied space in cache storage 105 will be made available for storing object *x*, and that may be wholly or partially overwritten by object *x* when it is stored in cache storage 105. In other words, a separate overwriting step to first erase object *y* is not required before object *x* is stored in cache storage 105. Once object *y* is evicted, object *x* is stored. File 120 is also updated to reflect eviction of object *y*, storage of object *x*, and any change in average object size. In the described embodiment, object *y* is evicted based on a

Least Recently Used (LRU) basis; the date and/or time of the last request for object *y* may be determined from file 120. However, other eviction algorithms could be used. Moreover, it may not be necessary to evict objects if the cache has room to store object *x*.

[0024] FIG. 6 is a flow chart showing operation of an embodiment of the invention. Beginning at step 412, the web cache receives a request for an object. At step 404, a determination is made as to whether the object is currently cached. If so, the request is serviced from the cache at step 406, the appropriate data records updated (e.g., date/time of request) at step 408, and the process ends. If it is determined at step 404 that the requested object is not cached, the object is retrieved at step 412, and the request serviced at step 414 (i.e., the object is served in response to the request). At step 416, a determination is made as to whether the object has previously been requested. If it is not the first request, flow proceeds to step 418. At step 418, one or more objects are evicted from the cache if necessary to make room for the new object, the new object is stored, and the appropriate data records updated. After step 418, flow proceeds to the end. If this is the first request for the object, flow would branch at step 416 to step 422. At step 422, a determination is made as to whether the size of the retrieved object is less than the average size of currently cached objects. If yes, the retrieved object will be cached, and flow proceeds to step 418. If the retrieved object is larger than the average size of currently cached objects, the appropriate data records are updated at step 424 (e.g., date and/or time of request), but the retrieved object is not cached, and flow proceeds to end.

[0025] Selective web caching according to the present invention was tested in a simulation using web access traces collected from a major Internet node, and which contained more than 13 million cacheable requests during an eighteen day period. In the simulation, cache size was varied from 17% to approximately 11% of total traffic volumes. Hit Ratio for selective web caching (SWC) according to the invention was compared to Hit Ratio for web caching using only LRU eviction (LRU-only caching). Also compared was the number of insertion and eviction operations per request for SWC versus LRU-only caching. FIG. 7 depicts Hit Ratios, at different cache sizes, of SWC and of LRU-only caching. As shown in FIG. 7, SWC consistently achieved Hit Ratios throughout the range of simulated cache sizes that were higher than the Hit Ratios for LRU-only caching. SWC improvement to Hit Ratio ranged from 1.8% to 5.5%.

[0026] FIGS. 8 and 9 illustrate the average number of insertion and eviction operations per request for SWC and LRU-only caching. SWC shows significantly lower insertion/eviction operations across the entire simulated cache size range. Notably, SWC reduced insertion operations by 30% to 40%, making higher cache throughput possible.

[0027] FIG. 10 illustrates the total number of disk operations at different cache sizes for SWC and LRU-only

caching. Compared to one operation per request for LRU-only caching (see above), SWC operations per request range from 0.77 to 0.81. Notably, this is achieved in conjunction with consistently higher Hit Ratios.

[0028] Although specific examples of carrying out the invention have been described, those skilled in the art will appreciate that there are numerous variations and permutations of the above described systems and methods that fall within the scope of the invention as set forth in the appended claims. For example, a machine-readable medium could have machine-executable instructions stored thereon such that, when the instructions are read and executed by an appropriate device (or devices), steps of a method according to the invention are performed. As another example, eviction schemes other than LRU could be employed instead of, or in conjunction with, LRU eviction. As yet another example, various parameters in the above described methods and systems could be varied. Instead of caching an object on a second request, the object could be cached on the third or other subsequent request. Instead of caching an object if its size is less than the average of the currently cached objects, the object might be cached if its size is less than a median cached object size, a percentage or other multiple of the average cached object size, or using some other size criterion- The invention might be implemented in environments other than web caches. The invention might also be implemented in environments where a request for an object is generated internally, either in response to receipt of an externally-generated request, or independent of an external request. These and other modifications are within the scope of the invention as defined by the attached claims.

Claims

1. A method of managing a computer memory by selectively admitting new objects into the memory, comprising:

receiving a request for an object not stored in the memory (412);
 storing the object in the memory if it has previously been requested a designated number of times (416) or if the size of the object satisfies a size criterion (422); and
 evicting, as may be necessary in connection with storing the requested object, one or more currently stored objects.

2. The method of claim 1, further comprising:

determining if the object has previously been requested the designated number of times; and
 comparing, subsequent to the determining and only if the object has not previously been requested the designated number of times, the ob-

ject size to the size criterion.

3. The method of any preceding claim, wherein:

storing the object in the memory if it has previously been requested comprises storing the object in memory if the object has been requested one prior time, and
 storing the object in memory if its size satisfies the size criterion comprises storing the object in memory if it is smaller than the average size of objects currently scored.

4. The method of any preceding claim, further comprising:

updating a first data record to indicate that the object has been requested.

5. The method of claim 4, further comprising:

updating a second data record to indicate presence of the object in the memory; and
 recalculating the size criterion based upon the presence of the object in the memory.

6. The method of any preceding claim, wherein the one or more currently stored objects are evicted on a Least Recently Used (LRU) basis.

7. The method of any preceding claim, wherein receiving a request for an object comprises receiving a request for an object having a Universal Resource Locator (URL).

8. The method of claim 7, wherein receiving a request for an object having a Universal Resource Locator (URL) comprises receiving a request for web content available from an origin web server.

9. A server, comprising:

at least one network connection permitting communication with client devices and with one or more origin servers providing objects;
 a memory (105); and
 a processor (110) configured to control admission of objects into the memory by:

receiving a request for an object not in stored in the memory and retrieving the object from an origin server;
 storing the object in the memory if it has previously been requested a designated number of times or if the size of the object satisfies a size criterion; and
 evicting, as may be necessary in connection with storing the requested object, one or

more currently stored objects.

10. The server of claim 9, wherein the processor (110) is configured to determine the object has previously been requested the designated number of times; and compare, subsequent to the determining and only if the object has not previously been requested the designated number of times, the object size to the size criterion. 5
11. The server of claim 9 or 10, wherein the processor is configured to store the object in the memory if it has previously been requested by storing the object if the object has been requested one prior time, and store the object in the memory if its size is satisfies the size criterion by storing the object if it is smaller than the average size of objects currently stored. 15
12. The server of any one of claims 9 to 11, wherein the processor is further configured to update a first data record to indicate that the object has been requested. 20
13. The server of claim 12, wherein the processor is further configured to update a second data record to indicate presence of the object in the memory, and to recalculate the size criterion based upon the presence of the object in the memory. 25
14. The server of any one of claims 9 to 13, wherein the processor is configured to evict one or more currently stored objects on a Least Recently Used (LRU) basis. 30
15. The server of any one of claims 9 to 14 wherein die server is a web server, wherein the requested object has a Universal Resource: Locator (URL), and wherein the processor is configured retrieve the object from an origin server based on the URL. 35
16. The server of claim 15, wherein the requested object is web content. 40
17. A computer readable medium having embodied thereon a computer program comprising program code means to perform the method as claimed in any one of claims 1 to 8. 45

Patentansprüche

1. Verfahren zum Verwalten eines Computerspeichers durch selektives Einlassen von neuen Objekten in den Speicher, das die folgenden Schritte aufweist: 50
 Empfangen einer Anforderung eines Objekts, das nicht im Speicher gespeichert ist (412);
 Speichern des Objekts im Speicher, wenn es

vorher eine ausgewiesene Anzahl von Malen (416) angefordert worden ist oder wenn die Größe des Objekts ein Größenkriterium erfüllt (422); und

Auswerfen, falls es in Verbindung mit dem Speichern des angeforderten Objekts notwendig ist, von einem oder mehreren gegenwärtig gespeicherten Objekten.

2. Verfahren nach Anspruch 1, das ferner Folgendes aufweist:

Feststellen, ob das Objekt vorher die ausgewiesene Anzahl von Malen angefordert worden ist; und

Vergleichen anschließend an das Feststellen und nur dann, wenn das Objekt nicht vorher die ausgewiesene Anzahl von Malen angefordert worden ist, der Objektgröße mit dem Größenkriterium.

3. Verfahren nach einem der vorhergehenden Ansprüche, wobei:

das Speichern des Objekts im Speicher, wenn es vorher angefordert worden ist, das Speichern des Objekts in Speicher aufweist, wenn das Objekt ein früheres Mal angefordert worden ist, und das Speichern des Objekts in Speicher, wenn seine Größe das Größenkriterium erfüllt, das Speichern des Objekts im Speicher aufweist, wenn es kleiner als die Durchschnittsgröße der gegenwärtig gespeicherten Objekte ist.

4. Verfahren nach einem der vorhergehenden Ansprüche, das ferner Folgendes aufweist:

Aktualisieren eines ersten Datensatzes, um anzuzeigen, dass das Objekt angefordert worden ist.

5. Verfahren nach Anspruch 4, das ferner Folgendes aufweist:

Aktualisieren eines zweiten Datensatzes, um das Vorhandensein des Objekts im Speicher anzuzeigen; und Neuberechnen des Größenkriteriums beruhend auf dem Vorhandensein des Objekts im Speicher.

6. Verfahren nach einem der vorhergehenden Ansprüche, wobei das eine oder die mehreren gegenwärtig gespeicherten Objekte auf einer "Least Recently Used"(LRU)-Grundlage ausgeworfen werden. 55

7. Verfahren nach einem der vorhergehenden Ansprüche, wobei das Empfangen einer Anforderung eines

Objekts das Empfangen einer Anforderung eines Objekts aufweist, das einen "Universal Resource Locator" (URL) aufweist.

8. Verfahren nach Anspruch 7, wobei das Empfangen einer Anforderung eines Objekts, das einen "Universal Resource Locator" (URL) aufweist, das Empfangen einer Anforderung eines Webinhalts aufweist, der von einem Ursprungswebserver abrufbar ist.

9. Server, der Folgendes aufweist:

mindestens eine Netzwerkverbindung, die eine Kommunikation mit Clientvorrichtungen und mit einem oder mehreren Ursprungsservern ermöglicht, die Objekte bereitstellen;
einen Speicher (105); und
einen Prozessor (110), der konfiguriert ist, den Einfluss von Objekten in den Speicher zu steuern,

durch:

Empfangen einer Anforderung eines Objekts, das nicht im Speicher gespeichert ist, und Abrufen des Objekts von einem Ursprungsserver;
Speichern des Objekts im Speicher, wenn es vorher eine ausgewiesene Anzahl von Malen angefordert worden ist oder wenn die Größe des Objekts ein Größenkriterium erfüllt; und
Hinauswerfen, falls es in Verbindung mit dem Speichern des angeforderten Objekts notwendig ist, von einem oder mehreren gegenwärtig gespeicherten Objekten.

10. Server nach Anspruch 9, wobei der Prozessor (110) konfiguriert ist, festzustellen, dass das Objekt vorher die ausgewiesene Anzahl von Malen angefordert worden ist; und
anschließend an das Feststellen und nur dann, wenn das Objekt nicht vorher die ausgewiesene Anzahl von Malen angefordert worden ist, die Objektgröße mit dem Größenkriterium zu vergleichen.

11. Server nach Anspruch 9 oder 10, wobei der Prozessor konfiguriert ist,
das Objekt im Speicher zu speichern, wenn es vorher angefordert worden ist, durch Speichern des Objekts, wenn das Objekt ein früheres Mal angefordert worden ist, und
das Objekt im Speicher zu speichern, wenn seine Größe das Größenkriterium erfüllt, durch Speichern des Objekts, wenn es kleiner als die Durchschnittsgröße der gegenwärtig gespeicherten Objekte ist.

12. Server nach einem der Ansprüche 9 bis 11, wobei der Prozessor ferner konfiguriert ist, einen ersten Datensatz zu aktualisieren, um anzuzeigen, dass

das Objekt angefordert worden ist.

13. Server nach Anspruch 12, wobei der Prozessor ferner konfiguriert ist, einen zweiten Datensatz zu aktualisieren, um das Vorhandensein des Objekts im Speicher anzuzeigen, und das Größenkriterium beruhend auf dem Vorhandensein des Objekts im Speicher neu zu berechnen.

14. Server nach einem der Ansprüche 9 bis 13, wobei der Prozessor konfiguriert ist, ein oder mehrere gegenwärtig gespeicherte Objekte auf einer "Least Recently Used" (LRU)-Grundlage auszuwerfen.

15. Server nach einem der Ansprüche 9 bis 14, wobei der Server ein Webserver ist, wobei das angeforderte Objekt einen "Universal Resource Locator" (URL) aufweist, und wobei der Prozessor konfiguriert ist, das Objekt von einem Ursprungsserver beruhend auf der URL abzurufen.

16. Server nach Anspruch 15, wobei das angeforderte Objekt ein Webinhalt ist.

17. Computerlebares Medium, auf dem ein Computerprogramm verkörpert ist, das ein Programcodemittel aufweist, um das Verfahren nach einem der Ansprüche 1 bis 8 auszuführen.

Revendications

1. Procédé de gestion d'une mémoire d'ordinateur en admettant de manière sélective de nouveaux objets dans la mémoire, consistant à :

recevoir une requête pour un objet non stocké dans la mémoire (412) ;
stocker l'objet dans la mémoire s'il a été demandé auparavant un nombre de fois désigné (416) ou si la taille de l'objet répond à un critère de taille (422) ; et
expulser, comme éventuellement nécessaire par rapport au stockage de l'objet demandé, un ou plusieurs objets stockés actuellement.

2. Procédé selon la revendication 1, consistant en outre à :

déterminer si l'objet a été demandé auparavant le nombre de fois désigné ; et
comparer, à la suite de la détermination et seulement si l'objet n'a pas été demandé auparavant le nombre de fois désigné, la taille de l'objet au critère de taille.

3. Procédé selon l'une quelconque des revendications précédentes, dans lequel :

- le stockage de l'objet dans la mémoire s'il a été demandé auparavant consiste à stocker l'objet dans une mémoire si l'objet a été demandé une précédente fois, et
le stockage de l'objet dans une mémoire si sa taille répond au critère de taille consiste à stocker l'objet dans une mémoire s'il est plus petit que la taille moyenne des objets stockés actuellement.
- 5
- 10
4. Procédé selon l'une quelconque des revendications précédentes, consistant en outre à :
- mettre à jour un premier enregistrement de données pour indiquer que l'objet a été demandé.
- 15
5. Procédé selon la revendication 4, consistant en outre à :
- mettre à jour un second enregistrement de données pour indiquer la présence de l'objet dans la mémoire ; et recalculer le critère de taille en se basant sur la présence de l'objet dans la mémoire.
- 20
- 25
6. Procédé selon l'une quelconque des revendications précédentes, dans lequel l'un ou les plusieurs objets stockés actuellement sont expulsés sur une base d'utilisation la moins récente (LRU).
- 30
7. Procédé selon l'une quelconque des revendications précédentes, dans lequel la réception d'une requête pour un objet consiste à recevoir une requête pour un objet ayant un localisateur de ressources universel (URL).
- 35
8. Procédé selon la revendication 7, dans lequel la réception d'une requête pour un objet ayant un localisateur de ressources universel (URL) consiste à recevoir une requête pour un contenu Web disponible à partir d'un serveur Web d'origine.
- 40
9. Serveur comprenant :
- au moins une connexion réseau permettant une communication avec des dispositifs clients et avec un ou plusieurs serveurs d'origine fournissant des objets ;
une mémoire (105) ; et
un processeur (110) configuré pour contrôler l'admission d'objets dans la mémoire en :
- 45
- 50
- recevant une requête pour un objet non stocké dans la mémoire et en récupérant l'objet à partir d'un serveur d'origine ;
stockant l'objet dans la mémoire s'il a été demandé auparavant un nombre de fois désigné ou si la taille de l'objet répond à un
- 55
- critère de taille ; et
expulsant, comme éventuellement nécessaire par rapport au stockage de l'objet demandé, un ou plusieurs objets stockés actuellement.
10. Serveur selon la revendication 9, dans lequel le processeur (110) est configuré pour déterminer que l'objet a été demandé auparavant le nombre de fois désigné ; et
comparer, à la suite de la détermination et seulement si l'objet n'a pas été demandé auparavant le nombre de fois désigné, la taille de l'objet au critère de taille.
11. Serveur selon la revendication 9 ou 10, dans lequel le processeur est configuré pour :
- stocker l'objet dans la mémoire s'il a été demandé auparavant en stockant l'objet si l'objet a été demandé une précédente fois, et
stocker l'objet dans la mémoire si sa taille répond au critère de taille en stockant l'objet s'il est plus petit que la taille moyenne des objets stockés actuellement.
12. Serveur selon l'une quelconque des revendications 9 à 11, dans lequel le processeur est en outre configuré pour mettre à jour un premier enregistrement de données pour indiquer que l'objet a été demandé.
13. Serveur selon la revendication 12, dans lequel le processeur est en outre configuré pour mettre à jour un second enregistrement de données pour indiquer la présence de l'objet dans la mémoire et pour recalculer le critère de taille en se basant sur la présence de l'objet dans la mémoire.
14. Serveur selon l'une quelconque des revendications 9 à 13, dans lequel le processeur est configuré pour expulser un ou plusieurs objets stockés actuellement sur une base d'utilisation la moins récente (LRU).
15. Serveur selon l'une quelconque des revendications 9 à 14, dans lequel le serveur est un serveur Web, dans lequel l'objet demandé comporte un localisateur de ressources universel (URL) et dans lequel le processeur est configuré pour récupérer l'objet à partir d'un serveur d'origine en se basant sur le localisateur URL.
16. Serveur selon la revendication 15, dans lequel l'objet demandé est un contenu Web.
17. Support lisible par ordinateur sur lequel est intégré un programme d'ordinateur comprenant un moyen de code de programme destiné à mettre en oeuvre le procédé selon l'une quelconque des revendications 1 à 8.




FIG. 1
Prior Art




FIG. 2
Prior Art




FIG. 3




FIG. 4




FIG. 5




FIG. 6




FIG. 7




FIG. 8




FIG. 9




FIG. 10

REFERENCES CITED IN THE DESCRIPTION

This list of references cited by the applicant is for the reader's convenience only. It does not form part of the European patent document. Even though great care has been taken in compiling the references, errors or omissions cannot be excluded and the EPO disclaims all liability in this regard.

Non-patent literature cited in the description

- A filtering algorithm for web caches. **KURCEWE-WICZ M et al.** Computer Networks and ISDN Systems. North Holland Publishing, 25 November 1998, vol. 30, 2203-2209 **[0011]**